Ray Brescia:

Welcome everybody. My name is Ray Brescia. I teach at Albany Law School. Some of you may have known me from a prior life as a lawyer at the Legal Aid Society of New York and at the Urban Justice Center. It is just my great privilege to be working on this session today with Adrián Palma from Microsoft. He's going to go through some wild technologies in real time and show you all how this generative AI stuff works. And then we're going to work together to do some interventions using generative AI. So without further ado, I'm going to turn over to Adrián, but actually with some further ado, Lisa Rivage, my colleague at Albany Law School, who is our CLE captain today, she's going to put her email in the chat. And if you need any of the affidavits or forms, please ask her for them, but you should have those in the emails that have been sent out on the session.

If you are online, and you are all online, but please, no cameras. And if you do have a question, raise your hand and I will get to you. Raise your hand throughout, you can ask a question at any time. And I think you're going to have lots of questions, your mind may very well be blown, but I'm going to turn it over to Adrián. He's got the overview of the session so you have a sense of what we're going to be talking about, and then he's going to dive right into it. So Adrián, take us away. And you're muted. So we got that out of the way first.

Adrián Palma:

There you go. Can you hear me?

Ray Brescia:

Yes.

Adrián Palma:

Perfect. Thank you so much, Ray, for that introduction. And it really has been wonderful to work with you and to explore the use of generative AI to scale access [inaudible 00:02:47] efforts. So I'm just going to quickly share my screen here so that you have a sense of the agenda for today's presentation. And so today we're really going to focus our on using low-code and no-code AI solutions to scale access-to-justice efforts. Just by way of background, I am Microsoft's pro bono program manager. Prior to joining Microsoft, I was managing a global law firm pro bono department at Akin Gump in Washington DC where I am headquartered. I am not an attorney, I am not a software engineer, but I strongly believe in the power of low-code, no-code technology to make us all global citizen developers and to be able to use our expertise in pro bono and your particular skillsets to be able to develop solutions that help your clients or help your organization.

So that really will be the focus of today's presentation, just helping you learn about some of these low-code, no-code tools. And we're going to start off by learning a little bit about Copilot Studio. I know that there was a question in the rapid-fire presentation about the use of Copilot Studio. So I hope that person is here in this presentation because we are going to delve into the low-code, no-code technology through Copilot Studio. And I will also be demoing a Copilot that I have been building with Microsoft colleagues to help scale access-to-justice efforts for dreamers or DACA recipients. I myself am a DACA recipient. I was born in Mexico, immigrated to the US at the age of six. And so this is a very personal issue for me. And then we'll go a little bit more into Power Automates and plugins and how they work with Copilot Studio

All of this may sound like a lot to understand right now, but I promise that the presentation that we have prepared for you will make it easy to consume this information. I will then demo a couple of AI prompts that Ray and I have been working on over the last couple of weeks for consumer debt lawsuit issues. And then we're hoping to build some AI prompts together with all of you through this presentation. And then we'll share some closing thoughts and offer some opportunity for questions and answers as well.

So just to give a little bit of an overview of Copilot Studio. Copilot Studio is a low-code, no-code platform created by Microsoft. It was announced in November of last year as a low-code, no-code tool to be able to build out your own Copilots or your own AI assistants. You're going to hear Copilot a lot. That is Microsoft's branding of our AI assistants. And so every single Microsoft product has a Copilot, Microsoft Word Copilot, Excel Copilot, Teams Copilot, et cetera. And through Copilot Studio, you can actually build out your own Copilots pretty easily. And just yesterday actually, Microsoft announced a new revamp to Copilot Studio where you're going to be able to just interact with a chatbot through natural language processing and be able to build out a Copilot really, really quickly and really easily via Copilot Studio. So if you haven't heard about this technology, if you haven't explored it, I highly that you explore some of it after this session. And of course you can always reach out to me if you have any questions as well. And I'm just going to give a very brief one-minute demo video on the DACA Copilot, and then we'll actually delve together into the elements that make up the DACA Copilot, but I want you to just have a high-level overview of this.

[Speaking on Demo Video]

How do you change the world?

So let me show you what we've built. The DACA Copilot starts off by introducing itself and its main functionalities. The user can either type directly into the Copilot, for example, "What is DACA?" and the Copilot will produce a generative AI response where you can click on the hyperlinks and verify the authoritative data sources that produced

that response. Alternatively, to promote accessibility, the user can use the Adaptive Cards to also produce a generative AI response with the same authoritative resources. Using the DACA Copilot, pro bono volunteers can also now learn about the monthly DACA clinics, skill up, record their time, and even sign up directly to volunteer for the upcoming DACA clinics. The DACA Copilot also allows a pro bono volunteer to automate the production of their clients' DACA renewal forms. The Copilots, by clicking a button, will initiate a local-intake chatbot on the user's computer. Once the user or the bono volunteer completes that intake chatbot, the DACA renewal forms will automatically be produced on that user's computer, thus guaranteeing the security and the privacy of the user's and the client's data. Now that we have to...

[End of Demo Video]

Adrián Palma:

That was just a very brief high-level overview of the DACA Copilot and some of the functionalities. And so now I'm going to delve a little bit more into these functionalities. And so the DACA Copilot has three main purposes. It serves as a DACA 101; it can answer legal and clinic questions pertaining to DACA. We've trained up the DACA Copilot with authoritative information directly from the USCIS website, the I-821D instructions I-765 instructions, et cetera, in order for us to reduce the possibility of hallucinations. And so as you saw in the video, when you type into the DACA Copilot, for example, "What is DACA?" It then produces a generative AI response for you. And it then does have those citations where you can go directly into the USCIS website or a specific form and verify the validity and the authenticity of that information.

And so we also at Microsoft promote Responsible AI Standards at every single opportunity that we can. And so through Copilot Studio, not only can you incorporate those generative AI capabilities, but as you can see on this third screenshot on the right, you can actually customize the way that the generative AI responses are being produced. So for example, I put the Microsoft Pro Bono logo, I put meta prompts, "Response generated by AI. Note that this response was generated by AI based on USCIS websites." So you can customize the way that the generative AI response and the output shows up to the user so that you can promote those Responsible AI Standards and let the user know and build that trust with the user that what they are seeing through this generative AI response has been vetted by Microsoft's Pro Bono Program, for example.

The DACA Copilot can also automate pro bono volunteer engagement and training as well. So let's start with the screenshot here on the left-hand side. This is just the introduction of the DACA Copilot. If I were to click the pro bono volunteer Adaptive Card, or if I were to type into the interface, "I want to sign up to become a pro bono volunteer," the user will get this sign-up form that you see here in the third screenshot.

So essentially this is a Microsoft form that lives within Copilot Studio. And as soon as you complete that form, a series of automations get sent to the user.

So the user will then be sent a training video, as you see on the fourth screenshot here, so that they can through the same user interface, train up and see all of the materials that they need in order to become a DACA volunteer. But then using Power Automate's low-code, no-code solutions, we also incorporated these automations for when that volunteer signs up, the clinic manager will receive an automated email, "Hey, Adrián Palma just signed up. He has an expected capacity date to volunteer on March 21st. He speaks Spanish. This data has now been added automatically to our SharePoint list to keep track of all the data." And then you as a pro bono volunteer who just signed up, you will get this email that you see here on the right-hand side that includes all of the relevant trainings and resources.

And so again, this is just one way that we can use Microsoft Copilot Studio to be able to automate a lot of the tedious processes that do take up a significant portion of an organization or a team's time. You can also automate the collection of volunteer feedback and the storytelling aspect of this. So if you, again, click that pro bono volunteer Adaptive Card, it then gives you the option to report feedback. And so this is where you could just complete the feedback as a pro bono volunteer, "Hey, I love the DACA clinic; these were my concerns." And you can automate the collection of all of that feedback. And of course you can automate the collection too of the storytelling piece, if you want to start to generate some marketing surrounding DACA and encourage other volunteers to sign up.

The really exciting part is the automation of the DACA renewal applications. And so this entire project came about from Microsoft's Global Hackathon. Microsoft hosts a global hackathon every September. They give us the opportunity as employees to take the week off so that we could develop solutions. And so this team, it was a group of about 20 Microsoft volunteers, got together, we decided, "Hey, there needs to be a better way to automate the completion of DACA forms in a way that's secure and that promotes the privacy of our clients." And so what we developed is essentially a Python script. When you run that Python script in your computer, it then initiates this intake chatbot here. And as soon as you complete that intake chatbot, it produces your USCIS forms.

And so through Copilot Studio, you can integrate those sort of automations via what's known as plugins. So here on the left-hand side, you see the enter space for the Copilot Studio. If I were to click automate DACA forms, it would trigger this local automation in my computer, it would start and run the Python script. I wouldn't need to do any of this. And you as a user would just see the first question, "What is your last name?" And as a user, you can start completing that intake chatbot. And as soon as you're done completing this intake-chatbot questionnaire, in your computer, you'll see the filled-out immigration forms. And so again, we chose this solution, because through Copilot

Studio, we can integrate that plugin that runs on a local user's computer. So all of the data that I am entering stays in my computer. The immigration forms that are automated stay in my computer.

So when we're thinking about using AI to scale access to justice, we need to think about that privacy element. And so this solution eventually will allow us to release a DACA Copilot that can be used by any dreamer nationwide. I mean, I've done my own DACA renewals five or six times at this point, so I don't need an attorney to review it, I just need a faster way to do this. And so through this solution and through incorporating the solution to Copilot Studio using a plugin, we can allow any dreamer, wherever they're located, whether it's a library, wherever they're located, they can launch this automation solution, fill out the questionnaire, and all of a sudden they have their DACA renewal forms automated. And so this is just an example of what the G-1145 looks like. This was our proof of concept when building out this solution during the hackathon itself. And so-

Ray Brescia:

So we have a couple of questions, Adrián, I'm going to jump in here, "Using this tool, is there still a risk of hallucination?" Number one. Number two, "How easy is it to update the forms or the law which guides the system?" And then, "Do you know of a similar type of program..." Or maybe this will be this September's hackathon, "... for asylum seekers?"

Adrián Palma:

So it is very easy through Copilot Studio, I'm just going to share my browser here, to ground the data for the Copilot Studio. And just as a disclaimer, Microsoft has turned off the GenAI features for this month, they're just going through a compliance-review process. So I can't actually demo these for you, but do want to show you how easy it is to just set up a Copilot and how easy it is for you to ground the Copilot just with your data. And so stepping back a little bit, actually, Copilot Studio uses a technique called RAG or Retrieval Augmented Generation. And essentially what that means is that you can tell the Copilot, "Hey, you're an LLM. Before you respond to the user, I want you to look at this data first. Don't respond to the user without going through this dataset," that you built in into the Copilots. And so that for us, for example, with the DACA Copilot, was just the USCIS authoritative resources. We're instructing it, "Only answer user output based on these websites and based on these PDFs."

And so that's RAG, or Retrieval Augmented Generation, in short. And it allows you to avoid the issue of hallucination by grounding the output of the Copilot with only the data sources that you specify. I will say though, that I have still noticed instances of hallucination. For example, for the DACA Copilot, I added generative AI capabilities to

talk about Microsoft's monthly DACA clinics. I asked if malpractice insurance is provided to volunteers, and it told me no, even though I had instructed it and had provided the malpractice insurance document to it. So that's the importance of user testing, so that you can ensure that if there are hallucinations, there's ways that you can remediate those.

And so through user testing, you can maybe even ask pro bono volunteers, "Hey, we have a DACA Copilot here, type in prompts, type in queries, and let us know if you see a hallucinations." And then at that point, we can go back into the Copilot and fine-tune it a little bit so that we avoid that specific query producing a hallucination. But I would say probably 95% of the time I have not seen any hallucinations in the Copilot itself. And so just very quickly, if you wanted to set up your own Copilot using Copilot Studio, you would log in to Copilot Studio. I believe a license costs 200 bucks. And then you would just click on Copilots, click on new Copilot, you put your Copilot's name. This is where you would then be able to add a website to your organization or to USCIS. And you hit create. And in a couple of minutes you have a Copilot that is only looking at that specific data source from that website that you added.

And once you do create the Copilots, let me just go back here, you do have the option then to add other GenAI functionalities. So here's that website that you can add, for example, to USCIS. Say you run a monthly DACA clinic and you already have a ton of information that's been vetted by your own attorneys, you can upload all of those files directly here. And again, within a couple of minutes you have a Copilot itself that is through RAG or through retrieval augmented generation only looking at that website and only looking at the files that you specified and producing the outputs related to that.

And then for the third question that you asked, Ray, our US immigration team at Microsoft is currently working on an asylum Copilot. They're actually using SharePoint Copilot to facilitate that. They're also experimenting with the use of Copilot Studio to build this out, but we're just in the process of seeing what mechanism works best to provide the most accurate results for the users. And at Microsoft, we're starting small. All of these Copilots are geared towards our internal users first, towards scaling our internal clinics. And once we then are able to develop the technology and vet it, at that point, we can then bring in external stakeholders or asylum seekers or DACA recipients themselves to use this technology. But before we do that, we have to go through an extensive responsible AI process. We have a responsible AI team at Microsoft that will go through redlining and red-taping and ensure that every single mechanism of this Copilot is built out in a responsible way.

So this is going to be a longer-term project, but in terms of just starting and experimenting, if you have a license to Copilot Studio, so I just shared with you, it's so easy to just be able to set one up quickly and then ground it with your data, and then

just start testing it out and seeing what outputs might be produced from there. And Ray, I completely forgot the second question that you asked. Did I answer it through what I-

Ray Brescia:

You did. So there's a follow-up, there've been lots of questions flying in. It's great.

Adrián Palma:

Keep asking them.

Ray Brescia:

So what about confidential information? You said before that the documents that are generated stay on your computer and, say on your system, but what about as you're sharing information, is it being shared with any entity?

Adrián Palma:

No. So the beauty of this is that none of the data that you're entering in it here, none of the prompts that you're entering into your Copilot Studio are used to train up any LLM. And so all of these products that Microsoft is releasing are built on top of Azure OpenAI. And so Azure takes that protection that has been provided to the world's largest enterprise companies, and adds layers of artificial intelligence by stacking Azure on top of OpenAI. So if you were to use, for example, Bing Copilot, and you're not using the enterprise version of it, that data would then get fed to OpenAI to train the LLM, or to Microsoft, for example. But this is where it's really important to ensure that you have the right enterprise licenses set up, because that's what then gives you the protection to be able to use this technology without the possibility of the data being sent to an LLM to be trained up. For example, here I just opened up Bing Copilot, you see the protected sign. So I could put in information here that's particular to my organization that might have confidential information related to Microsoft's Pro Bono Program, and none of that data is leaving my tenants. But if I did not have the protected version, that's where that data that you're inputting does get sent to refine and train up these LLMs.

Ray Brescia:

And so in reality, anything you might put into an email, assuming you're using Outlook, it's the same level of protection, right? If you're using-

Adrián Palma:

Yes.

Ray Brescia:

... Outlook, and you're using the Copilot as part of the Outlook enterprise, maybe I didn't use all those words correctly, but that's the idea, right? If you would-

Adrián Palma:

Yes.

Ray Brescia:

... put something confidential in an email to co-counsel or to your client, it goes through the same people it would as if you put it into Copilot, it's the same restrictions?

Adrián Palma:

Exactly. It's the enterprise protection offered by Azure.

Ray Brescia:

Okay. I think we've motored through a lot of the questions. Oh, what about foreign languages, languages other than English?

Adrián Palma:

Yeah. So as you can see here with the consumer law Copilot, there is a language functionality here. So if you wanted to start a Copilot in Spanish, when you set up the Copilot, you could specify the language there. Or if you set it up in English first and then you want to add additional languages, these are all of the languages that are offered. So you can easily just build up a Copilot that's in Spanish for Spanish-speaking DACA recipients, for example, or in any of these other languages listed here.

Ray Brescia:

Oh, "We need Haitian Creole." That's what one of the comments was, and I totally endorse that.

Adrián Palma:

Yeah [inaudible 00:22:52].

Ray Brescia:

Okay, so do you want to get to our workshop, are you ready for that, or what do you want to do next?

Adrián Palma:

So before we get to the workshop, I just want to go through the AI plugins a little bit more. And so a plugin, as you just saw through the document automation solution, is a

way that you can extend the capabilities of Copilots built on Copilot Studio. And so, one way that you could build these plugins is by using Power Automate's AI Builder models. And that's going to be really what this entire workshop is going to be focused on, jointly working together to build out these prompts that can then be integrated directly into Copilot Studio. And so here, for example, I'm showing you a prompt that I built out through AI Builder to generate a letter for a client who is looking to travel outside of the country using advanced parole.

So as you can see here through the prompt, and I'll just zoom in a little bit, it says, "Write a legal letter with the tone of an attorney for a DACA recipient looking to leave the United States by filing an advanced parole application. Use the following information provided by the client to draft the letter." And so you'll see those blue dynamic properties, that's the input that the user is then eventually going to add through that similar Adaptive Card that I showed when you sign up to become a volunteer. And then the last part, it says, "Ensure that you include relevant legal citations from DACA, including legal authority to travel outside the country." And so just for purposes of this workshop, I inputted those dynamic properties with some of my information. For example, I put the reason for traveling using advanced parole is because my father is sick, and I also have a conference to attend in Mexico, for example. And so within seconds, what was generated, and let me just share my screen to make-

Ray Brescia:

Yeah, and just so everyone knows, Adrián moves very quickly, and this seems like it's all overwhelming, but we are actually going to walk through step by step, and by the end of this session, you're going to be able to do something like this too.

Adrián Palma:

Exactly, exactly. And so this is just an overview of what's coming, but you see here on the screen, this is an AI-generated letter to justify my travel outside of the country using advanced parole. It's not perfect, every single AI product needs a human layer of review, but this gets you somewhere. You no longer have to start from scratch, and you could take this and rework it to your specific language or writing style, for example. But that's just an example of the way that you can easily build out these AI plugins for Copilot Studio to think about ways to automate a lot of these processes and then extend the power of Copilots. But before we get to the consumer debt law prompts, I do want to share here on the screen the first CLE code. I'll repeat it a couple of times. So number one, capital D, capital R, O-B-O-T, 1DRobot. I'll leave this up for five more seconds.

Great. So that was an overview of Copilot Studio and the ways that we've been using Copilot Studio specifically to build out a DACA Copilot. And so Ray and I over the last couple of weeks have been experimenting with Copilot Studio and with AI prompts and

thinking about ways that we can use AI for consumer debt law issues. And so I'm going to just demo a couple of prompts that we have built out before we then get into the workshop. But before we do that, prompt engineering's really important for you to know because the output of your AI solutions is only as strong as the prompt that you write. And there's been a lot of talk about prompt engineering, there're some of the hottest jobs in the market right now; OpenAI is paying like a million bucks to be a prompt engineer. That's a fad, it's going to fade, but it is still important for you to understand the fundamentals of a good prompt. And I say it's going to fade because these LLMs are only getting smarter, they're only getting better at understanding how we as humans speak and write, but for the time being, it is important for you to just have a sense of how to build a good prompt. And we'll drop some links into the chat later-

Ray Brescia:

I was just going to say that Adrián. So you don't have to be writing and taking notes feverishly, these links will be circulated in the chat, we'll circulate them afterwards. We're also going to be able to send around training videos for folks. And I'm going to volunteer Adrián, although he's really volunteered to do this, if we wanted to do a session in a few weeks where we walk through some of these initiatives step-by-step, he's willing to do that with us.

Adrián Palma:

Yes, yes. And I will show you some tricks as well to prompt engineering for you to not have to think about this formula, but the formula that Microsoft promotes for a good prompt is the goal, the context, the source, and the expectations. So you see here through this sentence how those elements fit into a specific prompt. So it's quite simple of a formula--your prompts can be very extensive, they can be hundreds of words--but this is the overall formula that would produce the best output for you. And so I encourage you to just keep this in mind. And as Ray mentioned, we will be sharing some links in the chat just so you could learn a little bit more about this. So we've gone a little bit into AI Prompt Builder via Power Automate, I showed you how the advanced parole letter engagement... Or sorry, the letter was generated. And so this is just what that user interface looks like in Power Automate, you can select various prompts. And so for purposes of this presentation, we're really going to be focused on the prompts that say, "Create text with GPT using a prompt."

Ray Brescia:

And where does one find Power Automate?

Adrián Palma:

Yes, very good question. So you'll go to office.com. And then once you open up office.com, you'll see a waffle up appear at the top, click on that waffle. And if your organization has a subscription to Power Automate, which most organizations do, you will then just click on Power Automate here, and then it will bring you up to Power Automate. And Power Automate, as background, is just a low-code, no-code solution to be able to automate a lot of your processes. So you saw that DACA training email that was automated. So I basically connected an automation, "When a form is completed, you will send out these emails." And it will include all of this information. That's what a Power Automate flow looks like. And again, it empowers all of us, even if you don't have technical skills, to be able to delve into the world of automation.

And so once you log into Power Automate, you're going to see the AI hub over here, click on that. And that's where you're then going to see the user interface that I previewed a little bit ago. It's just taking a little bit of time to load, but you'll see all of the various AI models. You can build documents, automation models, or even document-extraction models, but for purposes of this presentation, we're just going to focus on AI prompts. So now I'm just going to go through a couple of the prompts that Ray and I have built out just to give you a little bit of inspiration as to the power of this technology and how easy it is to build out these prompts. And then after this demo, we'll then dive into a workshop to build out prompts together jointly.

So let's start off with a prompt that's specific, for example, for generating an FAQ for your clients. And I'm going to open up another tool here called Microsoft Loop. This is where I have saved examples of complaints and summons particular to a consumer debt lawsuit. So I'm just going to take the text of that PDF, I just added it here. And so what I'm going to do is I'm going to drop that text here into this prompt, but before I do that, the prompt says, "Create an informative FAQ document based on the provided text..." And that's where I'm going to put the text of the complaint, "... addressing common questions, timelines, concerns, and any other relevant information. Please use simple language." And so let's test the prompt out. So here I'm going to drop in the text of the complaint. All of the information for the client has been changed, by the way, so there's no PII that's being divulged through this information. And then you click test prompt. And then as you can see here, now you have an entire FAQ document that you can send to your client that breaks down that complaint in a much easier way to understand, for someone who maybe this is their first time experiencing a legal issue, or maybe they just need a little bit of help understanding what they need to do in regards to a complaint or summons that they just received.

So that's one way that staff attorneys, pro bono coordinators, and even clients for example, can use these sort of tools to be able to drop the text of the complaint and then be able to generate these kind of documents. Legalese Decoder, this is another really cool prompt that we built out. As someone who is not an attorney, I sometimes

look at these statutes and I'm like, "What is this referring to?" So we created a prompt to help decode legalese. So the prompt says, "Define the meaning of the following legal or otherwise-hard-to-understand text." And that's where I'm going to drop a term, "Please use simple language and include a list of references to further understand this topic." So let's test the prompt out. I'm just going to drop CPLR Section 3212 directly here, and let's see what it produces for us.

Ray Brescia:

So the GPT is pulling from where?

Adrián Palma:

So this is pulling from the internet, so this is GPT-4 Turbo, to then find relevant information in the internet that can help out with this.

Ray Brescia:

But you could have cut and paste the actual provision in the law in there, like you did with the last-

Adrián Palma:

Mm-hmm.

Ray Brescia:

... prompt? Okay.

Adrián Palma:

Yes, yes, exactly. And so you see here the AI response gives you a breakdown that's a lot easier to understand than just seeing CPLR Section 3212. And it's just another way that you can use these prompts to help your clients or to help your team better understand some of these legalese terms that often make the law inaccessible for individuals. You can also, for example, produce a know-your-rights document. So this prompt says, "I want you to act as a knowledgeable presenter who can provide valuable information on knowing your rights when facing a legal issue type. In this presentation, you'll explain the legal rights individuals have in such situations, including the importance of understanding the debt collection process, responding to a lawsuit, and defending oneself in the court."

Just for purposes of time, I'm going to just run this prompt, but let me zoom in so individuals can read the prompt I'm [inaudible 00:34:04]. And so I'm just going to put consumer debt lawsuit here, and then I'm going to test the prompt. And so if your organization does a ton of work with building know-your-rights presentations, this is

just one way to start at least collecting some sort of information. You're no longer starting from a blank Word document, this could just be one way for you to at least have some sort of document that you can build off of to help educate your clients on their rights specific to debt collection lawsuit issues.

Ray Brescia:

And it's a rough first draft, it's not something you would just send directly to a client without reviewing?

Adrián Palma:

Exactly. And I should actually add that for all of these AI prompts that you're seeing here on the screen, because of Microsoft's Responsible AI Standards, a human needs to verify the output of all of these prompts. So if I were to send this specific prompt to a DACA client, for example, they could use the technology to input the text. Maybe they're trying to understand what advanced parole means, but I would get pinged first as a creator of this prompt. I would then get an email that tells me, "Hey, this is what AI produced. Do you approve of this text?" If you don't, you could write there on the spot review the email, edit it. Once you hit approved, that's then what gets sent to the client. So that's really important to emphasize, that all of these prompts that you build out have to go through that approval process. And it is just part of our Responsible AI Standards to ensure that there is human oversight as to every single AI-generated content that's being produced by these Microsoft tools.

Ray Brescia:

Well, that's your practice at Microsoft. If you're working with a client or someone in the community with whom you're working, one could imagine that this could be done by someone who is not represented, or an organization might not have the same human-in-the-loop standard. Not to say that we would advocate for that, but that's Microsoft's internal policy, which is a good one, but it doesn't apply to the Microsoft product that someone may have in the general public.

Adrián Palma:

That's correct. And so I think this makes us rethink the ways, maybe, that our staff, for example, may be handling responses that clients may be sending to us, or even intakes (for example, I'm going to show you here another referral prompt that we built out). And so maybe instead of having... When I was Akin Gump, I was answering phone calls from clients who wanted to know exactly which DC organization they could seek X help for X, Y, Z. Maybe the program coordinator role is going to evolve to be able to triage hundreds of requests very quickly. You're no longer typing out all of this information, you're just verifying the accuracy of it, you hit approve, and then it gets sent to the

client, for example. So I think we're in the middle of a transition as to what these jobs may look like as well, just given the potential for automation to help us expand, and help us scale the amount of help that we can give to individuals.

And so just very quickly, here's a pro bono referral prompt that we built out as well, "I want you to act as a resourceful AI assistant. Help me find at least 10 pro bono legal services organizations in X city that specialize in handling X pro bono matter type. Give me the org's name, mission, contact detail, website, etc." And so let's test this out. Let's put New York City and consumer debt lawsuit. And maybe-

Ray Brescia:

And this is how a couple of organizations run a helpline, that they already do this, where they've got the folks operating the helpline, have access to generative AI, and they curate the answers, and then the only person who has access to the answers is someone within the organization. And then they can communicate this out to the consumer who's looking for the assistance. So that's how you would envision using something this as well, correct?

Adrián Palma:

Exactly, it can amplify the impact of that hotline. Instead of only being able to help 10 clients with two hours, you can maybe help 100 clients with that same amount of time, because your role is then to just review these AI-generated responses, approve them, they get sent out to the client. So again, we're in the process of just a radical shift in the way that we can start to think about the ways that we can facilitate the delivery of pro bono legal services using these tools. And as you can see here, it gives me 10 organizations, their contact information, their websites, et cetera.

And then also, the last prompt that I'll share here is this one, Staff Attorney Email to Client with a Summary of Facts. So you've onboarded a client, you're going to help them with a consumer debt law issue, and you want to automate the production of an email that goes out to the client, that has all of the basic facts and information that the client needs. So the prompt says, "I want you to act as an assistant. I will provide you with the language cited in the complaint, and you will generate an email sent by staff attorney name, attorney title, at nonprofit organization, to the mentioned individual that includes a summary and necessary actions for them. And this email should be formatted using HTML and prettified just for purposes of putting it into an email quickly." And so let me just take that complaint again, I'm going to just paste that directly here, just put myself, I'm not an attorney, but just for purposes of this demo, and then I'll just put Microsoft, and then let's hit test prompt. And so this prompt can really help you just scale that work, that does take a ton of time, to send personalized emails.

Well, I think the HTML addition. I'm just going to remove that. The HTML addition may have impacted it, but I did take a screenshot in case we had issues with demos. And so this is at the response from that prompt. So it's like, "Dear Maria Rodriguez..." Which is a fake name in that complaint, "... hope this email finds you well." Then it gives the client's basic summary of the complaint, and it also includes just information about the Pro Bono Program, how we can help them. And then it's signed by you as a Pro Bono Program manager at Microsoft. So yeah, that's again just another way that you can take some of the workflows and that may take up a significant amount of your time, use AI to produce that automated email. You then just have to approve it or revise it. You hit approve, it gets sent to the client. And you can focus on helping other people with that time that you saved.

Ray Brescia:

All right, so let's do one of these letters. Do you want to do the identity theft letter?

Adrián Palma:

Yes. And so actually, Ray and I created a template for this prompt, but it would be really great for the experts who are attending this breakout session to give us their input,. Maybe there's other elements that you'd like to see in this letter here. So I'm just going to read the prompt. And then, Ray, I can't see the participants or you, so just speak out loud if there's individuals who are putting into the chats potential revisions to this. So the prompt says, "Write a letter to, creditor name, at, creditor address, on behalf of, client name, who lives in, client address, and who has experienced issues such as identity theft, and who has just received a letter from the same creditor seeking to collect on an alleged debt. Include language regarding the filing of report with the Federal Trade Commission," assuming that the client has already done that. Before we test the prompt out, are there any other-

Ray Brescia:

No, no, no. Let's run this prompt, and then we'll come back to it and we'll tweak it as people have changes.

Adrián Palma:

Awesome. So these are, Massachusetts Avenue, Northwest. I'm just typing in anything. I live in DC, so you're going to see... Jane Doe, 5 Penn Avenue, Northwest, DC, 2009. Identity theft is the issue that we've identified here. We're going to hit test prompt. And here you go, you have a letter address, "Subject: Dispute of Alleged Debt and Identity theft. Dear creditor's name, I'm writing to you on behalf of my client who resides at... We recently received a letter from your company regarding an alleged debt that my client is being asked to pay. However, I would like to bring to your attention that my client has

been a victim of identity theft, which has resulted in fraudulent accounts being opened in their name." And then here's additional information about identity theft.

"My client has taken immediate action by reporting the incident to the appropriate authorities, including filing a report with the FTC. The FTC is responsible for investigating. Given the circumstances, it is crucial that we resolve this matter promptly and ensure that my client is not held responsible for any debts incurred fraudulently, and kindly request that you suspend any collection activities related to this debt until a thorough investigation can be conducted to determine the validity of the debt." And then "I have enclosed copies of the relevant documents, including the FTC report and any other supporting evidence to assist you in understanding the gravity of the situation. I urge you to review this information carefully. Please provide us with a written response within [a number of] days of receiving this letter, acknowledging the dispute, and confirming that all collection activities will be suspended until the investigation is complete. Additionally, I request that you provide us with any documentation or evidence you may have regarding the alleged debt. We appreciate your cooperation. And then thank you for your attention to this matter. Sincerely, name." As you can see here, for some reason the-

Ray Brescia:

It's not pulling them in. Yeah.

Adrián Palma:

It didn't pull them in, but this gives you a sense of still the power of AI, it would've pulled in the client's name, the address, all of that information here. So yeah, we'll stop there. I mean, what are some initial thoughts of this AI-produced identity theft defense letter, are there elements missing that we could potentially incorporate into the prompt?

Ray Brescia:

You can raise your hand, you can put it in the chat, whatever you want to do. Do you want to go back to the prompt, Adrián?

Adrián Palma:

Yeah, and I'll zoom in just so-

Ray Brescia:

So I'll start us off. You could include language regarding the filing of a report as optional, right?

Adrián Palma:

Yes, if the client has filed one. And then this is where you would put a dynamic value, and this could be yes or no.

Ray Brescia:

So it's almost like that last sentence is bracketed connected to that dynamic value?

Adrián Palma:

Exactly. The dynamic value serves as a condition. So I'll include that language if the client has filed one, depending on what you enter, so either yes or no.

Ray Brescia:

Okay. What about adding whether or not the client may be judgment-proof regardless of the question of the identity theft? This would be a second issue that could be added here.

Adrián Palma:

Sorry, Ray, I just closed this out because it looks like I was just having some issues with the prompt, so-

Ray Brescia:

No worries.

Adrián Palma:

... restarting this. Okay. Can you-

Ray Brescia:

So the question would be, I would want to say, "In addition my client is also judgment-proof (yes, no,. so please refrain from any further action should you find that this is not an instance of identity theft," something like that. And I'm going to rely on the consumer advocates to tweak that if there's better language. We've also got a question, "Is there a value to ask for multiple drafts so the user can pick the one they like best to start with?"

Adrián Palma:

So that's actually a really good question, because the thing with AI work products is that they are non-deterministic. And by that I mean that every single output will look a little bit different. So you could rerun this prompt and you'll see that there will be tweaks to the letter. If you are looking for specific language to include, we have to do the work that we're doing right here, to tweak the prompt, and make sure that there are dynamic

values that allow us to include that language if we want. But just as practice, if you wanted to rerun and test the prompt again, you will see that every single instance will produce a bit of a different letter for you. And so that's one way that you can explore different letters that may work best. But again, it really comes down to the prompt. So if you're looking for elements that are missing, that's what needs to be tweaked in the prompt so that you can see that output.

Ray Brescia:

So you couldn't say, "I like this letter," or, "I've tweaked this letter. And I love all the language in it, I've reviewed it very carefully, I've made changes to it, and this is the letter that I want to be generated, just with these changed values," you couldn't do that?

Adrián Palma:

So you can actually. So for example-

Ray Brescia:

We also had another question about requesting all proof of the debt the creditor has in all documents the creditor plans to use if the case proceeds to trial.

Adrián Palma:

Yes, yes. And so if you want that information to be contained every single time that you run this prompt, you need to include that as a sentence in the prompt. I mean, the other thing that I'll mention here is that you can actually take this generated AI response and then be able to use other Copilots. So say for example, I'm using the protected version of Bing Copilot, I'm using Microsoft Edge right now, I just click this little button, and all of a sudden I have Bing Copilot, and I could just copy and paste that AI-generated response, and I could put it into Bing Copilot and say, "Hey, can you rewrite this in a way that's more formal, more professional," or "can you write this letter in Spanish for a Spanish-speaking client?" For example.

So the AI prompts can get you a work product, and then you can start stacking Copilots on top of each other, to take the strengths of Bing Copilot, which language interpretation is one of the biggest strengths, and then be able to take this AI-generated identity theft defense letter, copy and paste that, drop it into Bing Copilot it, and then do other things to the letter. Or as the human in the conversation here, if the letter looks 90% okay, it's still missing a couple of sentences, I mean, that's where you can easily just write out those sentences as well. None of these AI work products are ever going to be perfect. And again, there always needs to be a human element to it, but those are two strategies that you can employ to get the AI work product through these AI Builder prompts, and then just either use your own voice and your expertise, or

use other Copilots, like Word Copilot, Bing Copilot, to then further refine that letter, for example.

Ray Brescia:

So all the pistons are firing. So I'm going to summarize a couple of the questions. Is there a way to integrate an existing template that the organization has? Can you train the assistant in the templates that an organization already has? Let's put it that way.

Adrián Palma:

Yes, yes. And that actually goes back to the question you had also just asked right now, Ray. And so I'm just adding a sentence here at the end of the prompt, "Write the letter in the format of this text." And that's where if you already have a template, if you already have existing language, that's where you can put that text directly into this. So it's going to generate the letter and match it as close as possible to the template text that you add directly into this. So that's one way that you can maintain your organization's voice in the production of these AI letters. You can instruct the prompt to generate it, but make sure that it's grounded by the voice and the style of the templated text that you drop into the prompt too.

Ray Brescia:

Could you give it the template and say, "Use this template and write a letter to Discover Credit on behalf of my client Jane Doe, instead of Chase Bank or Mary Smith," and that way it would sort of trick the AI to use the template and just change the key variables here?

Adrián Palma:

You can. And I actually built one out similar. It's one of these, it might be Answer to The Summons or The Complaint. Yeah, so this is where I put, for example, "Draft a legal answer in response to the summon /complaint for defendant, defendant name. Ensure that the legal answer follows the same style, tone and language as a simple summons answer." So that's where you can use these dynamic properties to change up the name of the defendant, the name of the plaintiff, for example, and then you can drop the summons answer or the complaint answer text that you may already have in your organization and instruct the prompt to use that as the format and the style when drafting this new response.

Ray Brescia:

Because I can hear the vibrations through the Zoom, that folks are thinking, "Well, if I've got to review carefully every letter that this generates to make sure that it hasn't

changed things, it hasn't introduced new ideas that are irrelevant, I'm just going to go back to the way I usually work, which is to call up the letter I sent to Chase last week, and now I'm just going to change my client's name and put in Discover. That's going to be a lot easier than having to proofread a new letter every time that the generator generates one."

Adrián Palma:

Yes. And two reactions to that. GenAI is not the solution to everything. For this particular instance, I think a templated letter actually works much better. You already have the language, and maybe you have highlights that need to be replaced for every new client. I think that's a way better infrastructure than an AI prompt. The other piece to this too, again, really exploring all of the Copilots that exist. I mean, within Word Copilot, for example, this is what Copilot looks like within Word, you could actually tell it, "Write me a letter-

Ray Brescia:

Yeah, you're doing exactly what someone's just commenting one could do.

Adrián Palma:

Yeah, "So write me a letter for X client using..." And this is where you can then reference a file. So if you have a template, you can then just put that file right there. And so Word Copilot is understanding your template, and then it's going to write you a letter for that client using that template. So that, I think would be the best applicability of GenAI if you are looking to use GenAI for that use case. But again, GenAI doesn't have to be applicable for every use case. I think the template works perfectly fine, and that's what I would personally actually choose for this particular use case.

Ray Brescia:

All right, do you want to go back to... Oh, you want to run this and see how it works?

Adrián Palma:

Sure. Let's see.

Ray Brescia:

Do we have all the values in there?

Adrián Palma:

So we don't have a sample-

Ray Brescia:

Answer.

Adrián Palma:

... answer.

Ray Brescia:

Okay, so the prompt, you've added these dynamic values simply by clicking on add dynamic value. And then underneath the prompt it then shows you what all the dynamic values are and you plug them in.

Adrián Palma:

Exactly. And that is low-code, no-code technology. In the past, without low-code, no-code, you would have to code all of this, right?

Ray Brescia:

Yeah.

Adrián Palma:

You would have to code dynamic values and how it relates to the prompt. So that's the beauty of these dynamic values, you just click a button. And I also do want to point this last sentence out. I was thinking, when generating this response to a complaint, prior to generating that response to the complaint, maybe you had a call with your client, and maybe you had a call on Zoom or on Teams, you recorded the call, you have a transcript, the client gave you information, they talked about the particular circumstances. I mean, you could even download that transcript from that call and then use the client-information dynamic value to paste that transcript. And so it'll take all of that information that the client provided to you and use that to also build out the response to the complaint. So that's just another way that, again, just thinking about the various uses of technology, this is GenAI, but a transcript is super simple technology, that you just hit record and you have that transcript information. That's another way that you could even further supplement this template letter or this GenAI-produced letter in response to the complaint based on particular client facts that were shared with you during a call or an intake, for example.

Ray Brescia:

Okay, so can you put up the second CLE code please?

Adrián Palma:

Yes. So the second CLE code is right here on the screen, 1DTool.

Ray Brescia:

And I understand you have to read it twice. So 1DTool.

Adrián Palma:

1DTool.

Ray Brescia:

Yeah. So can we go now to some of the mobilization things that you were thinking of?

Adrián Palma:

Yes. So Ray and I were also working on a mobilization campaign prompt. So say you wanted to build up an email marketing strategy to be able to reach more individuals, say in Brooklyn, that may not have knowledge about your organization, you can use these sort of prompts to generate that campaign for you. So this prompt says, "Design an email campaign for time period, for X pro bono service targeting the following X population, providing a bulleted list which contains a week, goal and sample email copy." And so let's just test this prompt out. So let's say we want to do a one-month email marketing campaign; this is going to be for consumer debt lawsuits. And then we could choose population, low income, New York City [inaudible 00:58:41]. Test the prompt. Sometimes it's really fast, sometimes it's really slow, but something went wrong. And that is why we have screenshots, if that happens with a... Sorry, let me just zoom out of here. Sorry.

So this is that campaign, the zoom is not working well. Okay, so here's a one-month email campaign targeting low-income New Yorkers for a pro bono service related to consumer debt lawsuits. Week one, go introduce a pro bono service and create awareness." And here's just a sample email copy. You can, again, tweak this, make it your own. "Week two, provide educational resources on consumer debt and legal rights." And then additional sample email copies. And then you could even expand the prop further. Say you're planning a rally to advocate for a specific bill that's going to expand protections for consumers. You could expand the prompt to say like, "Also give me a list of 30 slogans that we can use. And then give me a list of five trusted vendors in New York City that can help produce posters for us." So you can take this simple prompt and then just expand it to your own organization's needs to be able to build out these mobilization campaigns pretty easily and effectively.

The other prompt that I want to point us to, let me see, where is it? It is a prompt for language accessibility. So if you have a client that doesn't speak English, for example. I'm just going to share the prompts here. Give me one second. Okay. So this was a

prompt, Language Accessibility for Clients, "Act as a professional translator and translate X text into X language." So this would be really helpful for a Spanish-speaking client, for example, that received a summons, they don't know what to do with it, or they received a complaint. You can use this prompt to be able to translate the entirety of the complaint for your clients. And then you can further that prompt by saying, "Then provide me a bullet-point list of all of the key facts in the complaint," in that specific language that you asked for it. So not only will it translate the actual document for you, but then it can provide you with almost like bullet points. And I have an example of that here. So these are just bullet points for your clients in terms of what they need to do with next steps, in terms of maybe your own staff attorneys, this might be helpful, for next steps for them as well. But just wanted to also just point out that these prompts could be really powerful for language accessibility. And I know that's a theme that's been discussed throughout today's conference.

Ray Brescia:

So there was a request to show the code again, if you can, Adrián.

Adrián Palma:

The-

Ray Brescia:

The CLE code.

Adrián Palma:

Oh.

Ray Brescia:

Not the code code, the CLE code.

Adrián Palma:

Yeah. You made me start to sweat a little bit.

Ray Brescia:

Yeah, yeah, yeah, reveal your code.

Adrián Palma:

Yeah. So here you go, it is 1DTool.

Ray Brescia:

Okay, let's open the floor for questions, comments.

Adrián Palma:

I know a lot of information was thrown out, so feel free to ask any question. Let's see.

Ray Brescia:

I see there's stuff in the Q&A, but we've talked about those Q&A questions.

Adrián Palma:

I wonder if the audience may want to just build out a prompt together with the time-

Ray Brescia:

We have a question. Zoe, go right ahead.

Zoe:

Okay, thanks. Can you hear me?

Ray Brescia:

Yep.

Adrián Palma:

Yes.

Zoe:

Okay, cool. Adrián, I had put this in the chat, it's a long one, but I was curious about the instructions on guard rails, I guess we can call them, that Copilot can follow, because the GPTs that I've created in ChatGPT with respect to following instructions after probably at least one or two dozen prompts on average. And I was also curious, because Copilot is the only AI I've ever interacted with that spontaneously closed a conversation with it. I was challenging it basically. I was said, "Actually, I don't really buy what you're saying." I was challenging it. And it said, "I'm not going to talk to you anymore." So I get those two things, really.

Adrián Palma:

Yes. So Copilot, I mean, those are the guard rails that Microsoft has built in to prevent the technology from producing information that could be harmful. And we saw this initially when Bing Copilot came out to the world. There was a New York Times reporter who started talking to Bing Copilot, and then they were pressing it with questions. And Bing Copilot apparently had ulterior motives and a different personality. And so

Microsoft through the Responsible AI Standards has put guard rails to prevent that. And I wonder if the way that you referenced it may have triggered that. And so I don't think that there's an easy solution to it other than reporting it through feedback back to Microsoft potentially, but that experience happened because of the guard rails that they're trying to set up so that there are no prompt injections, for example, that happen to these Copilots.

And by prompt injections, I mean, for example, there's been almost jailbreaks of these Copilots where an experienced coder, or even just an experienced prompt engineer, will put a prompt in and it will make the Copilot spew information that it shouldn't have. So you're seeing just the Responsible AI Standards at play, but it does cause confusion. And so I think one way to remediate that, for example, would be through custom prompts that you build out for clients or for others in your organization, that are vetted, that have that responsible AI infrastructure built in through the approval process. I think you can prevent those instances from happening through that infrastructure, so that a client's experience isn't terrible or interrupted, as yours was, Zoe. And can you repeat the first question that you had?

Ray Brescia:

We'll come back to that. So we have another question from Mark O'Brien who asks, "Can you share access to a Copilot with volunteers who do not all work for the same organization?" I think you can, you do that now, right?

Adrián Palma:

You can. And so with Copilot Studio, once you've gone through the responsible AI infrastructure, once you've determined, "This is a Copilot that we want to share with others," you can publish Copilot Studio through various channels. And let me share my screen here again to give you a sense of that. So say I was ready to publish this, these are just some of the mediums in which you can publish. So you can create an external website, for example, and publish the Copilot directly there. Maybe you're working with another organization on Microsoft Teams, you can publish the Copilot into a Teams site and be able to use the functionalities for that sort of collaboration. This list is expanding here to include more third-party apps, but you can see Slack, for example, if that's something that your organization uses, Facebook. I've been trying to push for WhatsApp as well, just given that a lot of the clients I serve operate and speak to me primarily via WhatsApp. But yeah, once you have a Copilot set up, once you're comfortable with the output, you've done the user testing, gone through Responsible AI Standards and testing, then you can publish the Copilot through any of these mediums.

Ray Brescia:

Okay, we have a question. We have a little time left. John Zakour has a question.

Adrián Palma:

Maybe he's muted.

Ray Brescia:

Yeah.

John Zakour:

Oh, can you hear me now?

Ray Brescia:

Yep.

John Zakour:

Sorry. You keep referencing responsible AI. What is it?

Adrián Palma:

So responsible AI are a set of principles that Microsoft and other tech companies have started to develop and started to form cohesion to promote the responsible use of artificial intelligence. And so the approval mechanism that I keep mentioning throughout this presentation, that's one way to ensure that there's human oversight into all of the work product that AI is producing for you. So that is a responsible AI practice. But just more broadly, Microsoft has six tenets of responsible AI. I'll put a link to the tenets in the chats, but for example, one of them is fairness and transparency. And so that tenet essentially means when you're developing these AI chatbots or Copilots, you should think about the end user. And I think that's super important for our work. For the DACA Copilot, for example, it was a ton of dreamers at Microsoft that actually led the development of this. And so through our lived experience as DACA recipients, our mothers were the ones who essentially were the first point of contact to helping us with our initial DACA application. We've built all of that directly into this Copilot to promote as much accessibility and inclusivity for our DACA clients as possible.

And so responsible AI also encompasses that aspect, where the ideation, the creation, and the publication of these Copilots, at every single point, there needs to be an understanding of these six responsible AI tenets in order to create these Copilots and these AI systems that actually serve society and don't potentially cause harm. And at Microsoft, we have a responsible AI team. It's made up of dozens of individuals that, again, whenever you have a Copilot that you want to publish even internally, it has to go through an entire process of redlining and red-teaming and going through every single

interaction and output to ensure that nothing is going to hurt or impact the legal rights of our clients. And so responsible AI is essentially that mobilization, and the mobilization of also the principles to ensure that all of these AI systems and products are safe to use in society and do not cause material, legal or even life impact to our clients.

Ray Brescia:

Is there a way for folks to try some of this for free or on a limited basis?

Adrián Palma:

Yes, yes, there is. So Copilot Studio, you can sign up for a free 30-day trial. If you just want to experiment with Copilots, Bing Copilot is free, but because it is free, again, that data is getting sent back to train up the LLM. So I would not put any client information into the free version of Copilot. But think about some of your most tedious workflows, that's where you could start using Bing Copilot for free, to start experimenting with it. I'll give you an example, in this presentation, I used Bind Copilot to just produce me the entire run of show. And then I fed that into PowerPoint Copilot, and the PowerPoint that you saw today was all produced by AI. So I would encourage free experimentation through Bing Copilot, and then signing up for Copilot Studio. And my biggest piece of advice is YouTube. Just go to YouTube, train up on these technologies. Ray knows this, I send him tutorials all the time. That's how I learned about all these technologies as well. Just type in Microsoft Copilot Studio, and you have a ton of influencers who specialize in this technology and will teach you the main functionalities for free.

Ray Brescia:

We're certainly open to doing a hands-on workshop in the future after we get some material out to folks. I want to thank Adrián, this has been just so great. I don't know how you got so much information in 75 minutes, but it's truly been extraordinary. Please, there's a survey, which is a link in the chat, please fill out. And those of you who are interested in getting the CLE credit, please email those forms to my colleague Lisa Rivage at lriva@albanylaw.edu. And then the next session will begin in 10 minutes, we've got another little break. The next session will begin in 10 minutes, and you go to the original link to get back to that one. But join me in a round of applause for Adrián. That was great. Thank you so much. Thank you all for joining. And we look forward to generating more access to justice together. See you all again in 10 minutes. Thank you very much.

Adrián Palma:

Thank you all.